

UNITED STATES PATENT APPLICATION  
FOR  
METHOD AND APPARATUS FOR SHARING TLB ENTRIES

INVENTORS:

THOMAS E. WILLIS  
ACHMED R. ZAHIR

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP  
12400 WILSHIRE BOULEVARD  
SEVENTH FLOOR  
LOS ANGELES, CA 90025-1026

(408) 720-8300

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EL 672 751 252 US

Date of Deposit March 30, 2001

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231

Kristin Baker

(Typed or printed name of person mailing paper or fee)

Kristin Baker  
(Signature of person mailing paper or fee)

3/30/01  
Date

# METHOD AND APPARATUS FOR SHARING TLB ENTRIES

## FIELD OF THE INVENTION

[0001] This invention relates generally to the field of computer systems, and in particular to sharing translation lookaside buffer (TLB) entries among multiple logical processors.

## BACKGROUND OF THE INVENTION

[0002] Computing systems use a variety of techniques to improve performance and throughput. One technique is known in the art as multiprocessing. In multiprocessing, multiple processors perform tasks in parallel to increase throughput of the overall system.

[0003] A variation of multiprocessing is known in the art as multithreading. In multithreading, multiple logical processors, which may comprise a single physical processor or multiple physical processors, perform tasks concurrently. These tasks may or may not cooperate with each other or share common data. Multithreading may be useful for increasing throughput by permitting useful work to be performed during otherwise latent periods, in which the performance level of the overall system might suffer.

[0004] Another technique to improve performance and throughput is known in the art as pipelining. A pipelined processor performs a portion of one small task or

processor instruction in parallel with a portion of another small task or processor instruction. Since processor instructions commonly include similar sequences of component operations, pipelining has the effect of reducing the average duration required to complete an instruction by working on component operations of multiple instructions in parallel.

[0005] One such component operation, is a translation from virtual addresses to physical addresses. This operation is often performed by using a translation lookaside buffer (TLB). It is a function of the TLB to permit access to high-speed storage devices, often referred to as caches, by quickly translating a virtual address from a task, software process or thread of execution into a physical storage address. In systems which permit multiprocessing, including those systems that permit multithreading, identical virtual addresses from two different threads or software processes may translate into two different physical addresses.

[0006] On the other hand, multiple threads or software processes may share a common address space, in which case some identical virtual addresses may translate into identical physical addresses. To prevent mistakes in accessing high-speed storage, the data may be stored according to physical addresses instead of virtual addresses.

[0007] If a high-speed storage device is accessed by multiple logical processors, the size of the TLB may be increased to store virtual address translations for each logical processor or thread of execution. Unfortunately, the time required to

perform a virtual address translation increases with the size of the TLB, thereby reducing access speed and overall system performance. Alternatively, smaller faster TLBs may be physically duplicated for each logical processor, but physically duplicating these hardware structures may be expensive. Furthermore, in cases where multiple threads or software processes share a common address space, multiple TLB entries may contain duplicates of identical virtual address translations, thereby wasting these expensive resources.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings.

[0009] **Figure 1** illustrates a system level abstraction of a single processor.

[0010] **Figure 2** illustrates a dual processor system based on the system level abstraction of single processors.

[0011] **Figure 3** illustrates a dual processor system including a multiprocessor with shared resources.

[0012] **Figure 4a** illustrates one embodiment of a multiprocessor system with resource sharing.

[0013] **Figure 4b** illustrates an alternative embodiment of a multiprocessor system with resource sharing.

[0014] **Figure 5** illustrates one embodiment of a processor pipeline.

[0015] **Figure 6** illustrates one embodiment of a shared TLB used in an address translation stage.

[0016] **Figure 7** illustrates alternative embodiments of a shared TLB used in an address translation stage.

[0017] **Figure 8** illustrates one embodiment of control logic circuitry for use with a shared TLB.

[0018] Figure 9 illustrates alternative embodiments of a control logic process for TLB entry sharing.

[0019] Figure 10 illustrates one embodiment of a computing system including a multiprocessor with a shared TLB.

## DETAILED DESCRIPTION

[0020] These and other embodiments of the present invention may be realized in accordance with the following teachings and it should be evident that various modifications and changes may be made in the following teachings without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense and the invention measured only in terms of the claims.

[0021] Disclosed herein is a mechanism for sharing among multiple logical processors, a translation lookaside buffer (TLB) to translate virtual addresses, for example into physical addresses. The mechanism supports sharing of TLB entries among logical processors, which may access address spaces in common. The mechanism further supports private TLB entries among logical processors, which for example, may each access a different physical address through identical virtual addresses. The disclosed mechanism provides for installation and updating of TLB entries as private entries or as shared entries transparently, without requiring special operating system support or modifications. Through use of the disclosed sharing mechanism, fast and efficient virtual address translation is provided without requiring more expensive duplicate circuitry.

[0022] For the purpose of the following disclosure, a processor or logical processor may be considered to include, but is not limited to, a processing element

having access to an execution core for executing operations according to an architecturally defined or micro-architecturally defined instruction set. A processor or logical processor may at times, for the purpose of clarity, be logically identified with a machine state and a sequence of executable operations, also referred to herein as a thread of execution, task or process. The physical boundaries of multiple processors or logical processors may, accordingly, be permitted to overlap each other. For this reason, references may be made to a logical machine in order to distinguish it from a processor or logical processor, which may physically or functionally overlap with another processor or logical processor, these distinctions being made for the purpose of illustration rather than for the purpose of restriction.

[0023] Abstraction levels, such as system level abstractions, platform level abstractions and hardware level abstractions may, for the purpose of the following disclosure, be considered to include, but are not limited to, specified interfaces. Details of these specified interfaces are to permit design teams to engineer hardware, firmware or software components to work with, or communicate with, components of different or adjacent abstraction levels within a system. It will be appreciated that an implementation that supports or adheres to one or more of these abstraction level specifications further includes any necessary circuitry, state machines, memories, procedures or other functional components, the complexities of these components varying according to design tradeoffs. It will be further



appreciated that such variations and complexities are generally hidden by the associated abstraction level interfaces.

[0024] Figure 1 illustrates one embodiment of a system level abstraction of a single processor 110. Processor 110 includes a processing element, logical machine 111; a cache storage resource, L1 cache 112; a cache storage resource, L2 cache 113, and a data transmission resource 114.

[0025] Figure 2 illustrates a dual processor system 200 based on the system level abstraction of single processors from Figure 1. Dual processor system 200 comprises a central storage, memory 230; a first processor, processor 210 including logical machine 211, L1 cache 212, L2 cache 213, and data transmission resource 214; and a second processor, processor 220 including logical machine 221, L1 cache 222, L2 cache 223, and data transmission resource 224. It will be appreciated that not all of the logically identical resources need to be duplicated for each of the processors. For example, it may be more efficient to physically share a resource among multiple processors while preserving the logical appearance of multiple single processors, each having a complete set of resources.

[0026] Figure 3 illustrates a dual processor system including one embodiment of a multiprocessor 301 with shared resources, as part of a system 300. System 300 also includes memory 330. Multiprocessor 301 also includes first logical machine 311 having shared access to L1 cache 322 and a second logical machine 321 having shared access to L1 cache 322. Both logical machine 311 and logical machine 321

also have shared access to L2 cache 333, and data transmission resource 334. Shared L1 cache 322 and shared L2 cache 333 may be used, for example, to store copies of data or instructions transmitted via data transmission resource 334 from memory 330 for either logical machine 311 or logical machine 321.

[0027] Both logical machine 311 and logical machine 321 may access and exercise control over L1 cache 322, L2 cache 333 and data transmission resource 334, and so it may be advantageous to access data according to physical addresses for these shared resources to prevent mistakes. One way in which access and control may be provided to multiple logical machines, as shown in Figure 4a, includes a platform level abstraction (PLA) 411, and a hardware level abstraction (HLA) 414.

[0028] Figure 4a illustrates an embodiment of a multiprocessor 401 comprising a processor 410 that has access to exclusive resources 412 and shared resource 433 and also comprising a processor 420 that has access to exclusive resources 422 and shared resource 433. Resource 412 and resource 433 represent exclusive and shared resources respectively, for example cache resources, busses or other data transmission resources, virtual address translation resources, protocol resources, arithmetic unit resources, register resources or any other resources accessed through the hardware level abstraction 414. In one embodiment, access to resource 412 or to resource 433 is provided by the hardware level abstraction 414 through a corresponding mode specific register (MSR). For example, access to exclusive resource 412 is accomplished through hardware level abstraction 414 by providing

for PLA firmware to perform a write operation to the corresponding MSR 415.

Access to shared resource 433 is accomplished through hardware level abstraction 414 by providing for PLA firmware 411 to perform a write operation to the corresponding MSR 435. Sharing control 431 provides and coordinates access to shared resource 433 and to the corresponding MSR 435.

[0029] Similarly, access to exclusive resource 422 is provided through hardware level abstraction 424 by PLA firmware 421 performing a write operation to corresponding MSR 425. Access to shared resource 433 is provided through hardware level abstraction 424 by PLA firmware 421 performing a write operation to corresponding MSR 435 with sharing control 431 providing and coordinating access to the corresponding MSR 435, and thereby to shared resource 433.

[0030] Figure 4b illustrates an alternative embodiment of a multiprocessor 401 comprising a processor 410 and a processor 420 that have access to shared resources including register file 436, execution unit 437, allocation unit 438, and instruction queue 439. Additionally processor 410 has exclusive access to register renaming unit 416 and reorder buffer 417, and processor 420 has exclusive access to register renaming unit 426 and reorder buffer 427.

[0031] Instruction queue 439 contains instructions associated with a thread of execution for processor 410 and instructions associated with a thread of execution for processor 420. Allocation unit 438 allocates register resources from register file 436 to register renaming unit 416 for instructions in instruction queue 438

associated with the thread of execution for processor 410. Execution unit 437 executes instructions from instruction queue 438 associated with the thread of execution for processor 410 and then reorder buffer 417 retires the instructions in sequential order of the thread of execution for processor 410.

[0032] Allocation unit 438 further allocates register resources from register file 436 to register renaming unit 426 for instructions in instruction queue 438 associated with the thread of execution for processor 420. Execution unit 437 also executes instructions from instruction queue 438 associated with the thread of execution for processor 420 and then reorder buffer 427 retires the instructions in sequential order of the thread of execution for processor 420.

[0033] Modern processors are often heavily pipelined to increase operating frequencies and exploit parallelism. Figure 5 illustrates one embodiment of a processor pipeline wherein the front end of the pipeline includes instruction steering stage 501, address translation stage 502, and data fetch stage 503; and the back end of the pipeline culminates with instruction retirement stage 509. Data from successive stages may be stored or latched to provide inputs to the next pipeline stage.

[0034] The address translation stage 502 may perform a translation from a virtual address to a physical address using a storage structure called a translation lookaside buffer (TLB).

[0035] In one embodiment, an apparatus provides shared virtual address translation entries of a TLB 602 for use in address translation stage 502. Figure 6 shows a tag array 631 for storing virtual address data (VAD) which may comprise, for example, a virtual page number. The figure also shows a translation array 635 for storing: corresponding physical address data (PAD) which may comprise, for example, a physical page number; address space identifier data (ASID); attributes (ATRD) such as page size data, security data, privilege data, etc.; and other associated data. Tag array 631 includes data line 611 and corresponding sharing indication 616, data line 612 and corresponding sharing indication 617, other data lines and corresponding sharing indications and finally, data line 613 and corresponding sharing indication 618. Translation array 635 includes data line 621, data line 622, other data lines and finally, data line 623.

[0036] When data is read from tag array 631 and from corresponding translation array 635 it is may be latched by latch 633 and latch 637 respectively. Latch 633 includes both data portion 614 for storing virtual address data (VAD) and sharing indication 619 for identifying if the corresponding virtual address translation may be used in correspondence with a logical processor requesting the virtual address translation. The latch 637 includes, in data portion 624, a corresponding physical address data (PAD); an address space identifier data (ASID); attributes (ATRD) such as, page size data, security data, privilege data, etc.; and other associated data

for translating the virtual address and for checking if the latched output of translation array 635 may be shared.

[0037] Control logic 604 may use the data portion 614, sharing indication 619, and data portion 624 to identify if the virtual address translation is sharable. For example, if a processor initiates a TLB request to look up a virtual address translation and the TLB entry in latches 633 and 637 contains an ASID that matches the ASID for the virtual address to be translated, and further if the entry contains a VAD that matches the VAD for the virtual address, and finally if sharing indication 619 indicates a set of logical processes including one associated with the processor initiating the TLB request, then the entry in latch 633 and latch 637 may be used to translate the virtual address. Otherwise, control logic 604 may initiate installation of a new virtual address translation entry for TLB 602.

[0038] Whenever a miss occurs in TLB 602, the physical address data and other TLB data may be recovered from page tables in main memory. For one alternative embodiment control logic 604 may comprise a mechanism for recovering such data. Most modern processors use a mechanism called a page walker to access page tables in memory and compute physical addresses on TLB misses.

[0039] If a processor, either directly through software or indirectly through control logic 604, initiates a TLB request to installation of a new virtual address translation entry, the TLB 602 may be searched for any existing entries that can be shared. An entry retrieved from tag array 631 and translation array 635 may then

be latched by latch 633 and latch 637 respectively. If the TLB entry in latches 633 and 637 contains an ASID that matches the ASID for the virtual address to be translated, and further if the entry contains a VAD that matches the VAD for the virtual address, and finally if sharing indication 619 indicates a shared status, then the entry in latch 633 and latch 637 may be installed for the processor initiating the TLB request by adding the logical process associated with the initiating processor to the set of logical processes indicated by sharing indication 619 and thereafter the TLB entry may be used to translate the virtual address. Otherwise, control logic 604 may initiate allocation of a new virtual address translation entry for TLB 602.

[0040] If a processor, either directly through software or indirectly through control logic 604, initiates a TLB request to allocate a new virtual address translation entry, the TLB 602 may be searched for any invalid or replaceable entries. The retrieved TLB entry may then be reset by control logic 604 to contain an ASID that matches the ASID for the virtual address to be translated, a VAD that matches the VAD for the virtual address, a PAD that matches the PAD of the translated physical address, an ATRD that matches the ATRD of the translated physical address, and any other associated data corresponding to the virtual address translation. Finally the entry may be installed for the processor initiating the TLB allocation request by initializing the set of logical processes indicated by sharing indication 619 to contain only the logical process associated with the initiating processor. It will be appreciated that the sharing indication 619 may be

conveniently initialized by default to indicate a shared status for the virtual address translation. Alternatively if the allocation was initiated through software, for example, control logic 604 may initialize the sharing indication 619 by default to indicate a private status for the virtual address translation.

[0041] When it is desirable for a processor to purge a virtual address translation, the processor initiates a TLB request to look up the virtual address translation entry that translates the virtual address. The retrieved TLB entry may then be reset by control logic 604 by initializing the set of logical processes indicated by sharing indication 619 to the empty set. It will also be appreciated that the sharing indication 619 may be conveniently initialized by default to indicate a private status for the virtual address translation, for example, if no explicit invalid status is representable.

[0042] It will be appreciated that control unit 604 provides for efficient sharing of TLB 602 entries among logical processes without requiring additional support from, or modifications to, any particular operating system that may be selected for use in conjunction with a multiprocessor or multithreading processor employing the apparatus of Figure 6 to provide sharing of virtual address translations in an address translation stage 502. One such multiprocessor or multithreading processor may, for example, execute a 32-bit Intel Architecture (IA-32) instruction set which comprises IA-32 instructions of the Pentium® processor family. Another such multiprocessor or multithreading processor may, for example, execute a 64-bit



Intel Architecture (IA-64) instruction set which comprises IA-64 instructions of the Itanium™ processor family or may also execute a combination of both IA-32 and IA-64 instructions. Since such multiprocessors or multithreading processors may be used in various computer systems running any one of a number of operating systems, an apparatus employed by such multiprocessors or multithreading processors to provide sharing of TLB entries should accordingly be operating-system transparent, providing sharing of TLB entries among logical processes without requiring that the operating system actively manage the sharing of all TLB entries. It will also be appreciated that if a multiprocessor or multithreading processor has a mechanism to provide sharing of TLB entries in such a way that is operating-system transparent or operating-system independent, that it does not prohibit that multiprocessor or multithreading processor from also providing for additional operating-system support for managing some sharing of TLB entries.

[0043] Figure 7 illustrates alternative operating-system transparent embodiments of a shared TLB 702 used in an address translation stage 502. A scalable sharing indication scheme 703 comprises a status indication and a set of logical processes and associated processors for each corresponding virtual address translation entry in the shared TLB 702. Alternatively, the status indication may be implicitly represented by the set of logical processes and associated processors as illustrated in Figure 7b. As described above, control logic 704 may be used to

identify if a virtual address translation is sharable by the logical processors 710, 720, 740 and 780.

**[0044]** Shared TLB 702 stores virtual address translation entries 711 through 730.

A virtual address translation entry may include: a virtual address data (VAD) for example, a virtual page number; a corresponding physical address data (PAD) for example, a physical page number; an address space identifier data (ASID); attribute data (ATRD) such as, page size data, security data, privilege data, etc.; and other associated data for translating the virtual address and for checking if the virtual address translation entry may be shared. Each virtual address translation entry has, in shared TLB 702, a corresponding status indication (in Status 705) and a corresponding indication of the set of logical processes (in P 706) sharing the virtual address translation. When a processor requests a virtual address translation, TLB 702 will be searched for a valid virtual address translation entry having a VAD that matches the VAD of the virtual address to be translated. If the corresponding set of logical processes sharing the virtual address translation includes a process associated with the requesting processor, the entry retrieved may be used to translate the virtual address.

**[0045]** It will be appreciated that a set of logical processes sharing a virtual address translation may indicate inclusion of a process associated with a particular processor by simply indicating or listing that particular processor.

[0046] In Figure 7a, for example, a sharing indication corresponding to virtual address translation entry 711 indicates a private status of P and a set of logical processes of 0001, the low order bit being set to indicate that entry 711 may be used exclusively to translate virtual addresses for processor 710. Similarly a sharing indication corresponding to virtual address translation entry 713 indicates a private status of P and a set of logical processes of 0100, indicating that entry 713 may be used exclusively to translate virtual addresses for processor 740.

[0047] A sharing indication corresponding to virtual address translation entry 712 indicates a shared status of S and a set of logical processes of 0101, indicating that entry 712 may be shared and may be used to translate virtual addresses for processors 710 and 740. Similarly a sharing indication corresponding to virtual address translation entry 719 indicates a shared status of S and a set of logical processes of 1111, indicating that entry 719 may be shared and used to translate virtual addresses for all four processors 710-780.

[0048] A sharing indication corresponding to virtual address translation entry 716 indicates a invalid status of I and a set of logical processes of 0000 meaning that entry 716 may not be used to translate virtual addresses for any processor 710-780. It will be appreciated that the invalid status may be explicitly represented or implicitly represented by the corresponding set of logical processes. It will also be appreciated that one skilled in the art may produce other encodings to explicitly or implicitly represent sharing indications for TLB entries.

[0049] In Figure 7b, for example, a sharing indication corresponding to virtual address translation entry 711 may implicitly indicate a private status of P and an explicit set of logical processes of 01 meaning that entry 711 may be used to translate virtual addresses for processor 710. It will be appreciated that such an implicit status representation may permit any implicit private status to be changed to an implicit shared status if another processor is found that may make use of the corresponding virtual address translation entry.

[0050] For example, if a processor initiates a TLB request to look up a virtual address translation and the sharing indication corresponding to the retrieved TLB entry indicates a set of logical processes that does not include one associated with the processor initiating the TLB request, then the physical address data and other TLB data may be recovered from page tables in main memory. Control logic 704 may include a mechanism for recovering such data, or may invoke a mechanism such as a page walker to access page tables in memory and compute physical addresses. If the newly constructed virtual address translation matches the retrieved TLB entry, the requesting process may be added to the set of logical processes sharing the retrieved TLB entry. Otherwise the newly constructed virtual address translation may be installed in a new TLB entry for the requesting processor.

[0051] Figure 8 illustrates one embodiment of a control logic 804 for use with a shared TLB. Control logic 804 comprises storage cell 810, storage cell 811, and

storage cell 812. Storage cells 810 and 811 may be used to record set of logical processes sharing a virtual address translation entry. Processor  $P_0$  may be added to the set of logical processes sharing a virtual address translation by asserting the  $Share_0$  input signal to storage cell 810. Likewise, processor  $P_1$  may be added to the set of logical processes sharing a virtual address translation by asserting the  $Share_1$  input signal to storage cell 811. Either processor  $P_0$  or  $P_1$  may purge the translation by respectively asserting the  $Purge_0$  input signal to storage cell 810 or asserting the  $Purge_1$  input signal to storage cell 811. Storage cell 812 may be used to record a corresponding status for the virtual address translation entry. A shared status may be recorded by asserting the *Install Shared* input signal to storage cell 812. A private status may be recorded by asserting the *Install Private* input signal to storage cell 812.

[0052] Control logic 804 further comprises multiplexer 813 and OR gate 814. If a processor identifier (*PID*) for a logical processor requesting a virtual address translation is asserted at the select input of multiplexer 813, the output of multiplexer 813 will indicate whether the virtual address translation entry may be readily used to provide the virtual address translation for the requesting processor. If the set of logical processes indicates either logical processor  $P_0$  or  $P_1$  is sharing the translation then the output of OR gate 814 will indicate that the translation is valid.

[0053] It will be appreciated that modifications may be made in arrangement and detail by those skilled in the art without departing from the principles of the

invention disclosed and that additional elements, known in the art, may be further incorporated into control logic 804. It will also be appreciated that a control logic for operating-system transparent TLB entry sharing may comprise a combination of circuitry and also machine executable instructions for execution by one or more machines.

[0054] Figure 9a, for example, illustrates a diagram of one embodiment of a process for TLB entry sharing for a control logic 904. The process is performed by processing blocks that may comprise software or firmware operation codes executable by general purpose machines or by special purpose machines or by a combination of both. In processing block 910, a virtual address translation is accessed. In processing block 911, the sharability status of the virtual address translation is identified. In processing block 912, the result of processing block 911 is used to control processing flow. If a sharable status is identified, then processing flow continues in processing block 914, where a sharing indication with a shared status is provided.

[0055] Otherwise a private status is identified, and processing flow continues in processing block 913, where a sharing indication with a private status is provided.

[0056] Figure 9b illustrates a diagram of an alternative embodiment of a process for TLB entry sharing for control logic 904. In processing block 920, a virtual address translation is accessed. In processing block 921, the sharability status of the virtual address translation is identified. In processing block 922, the result of

processing block 921 is again used to control processing flow. If a sharable status is identified, then processing flow continues in processing block 927, where again a sharing indication with a shared status is provided. In processing block 928 a set of logical processes sharing the virtual address translation is provided.

[0057] Otherwise, in processing block 921, a private status has been identified, and processing flow continues in processing block 925, where a sharing indication with a private status is provided. In processing block 926 a logical processes using the virtual address translation is provided.

[0058] Figure 9c illustrates a diagram of another alternative embodiment of a process for TLB entry sharing for control logic 904. In processing block 930, virtual address translation VAT is accessed for processor  $P_i$ . In processing block 931, the sharability status of virtual address translation VAT is identified. In processing block 932, the set  $P_{VAT}$  of logical processes sharing virtual address translation VAT is checked to see if a process associated with processor  $P_i$  is indicated. The result is used to control processing flow. If processor  $P_i$  is indicated as sharing virtual address translation VAT then processing continues in processing block 938 where virtual address translation VAT is used to translate virtual addresses for processor  $P_i$ .

[0059] Otherwise, in processing block 932, processor  $P_i$  is not indicated as sharing virtual address translation VAT and processing continues in processing block 933, where a new virtual address translation  $VAT_i$  is built from page tables

and physical address data is computed for processor  $P_i$ . In processing block 934 the new virtual address translation  $VAT_i$  is checked to see if it matches the retrieved virtual address translation VAT. If so, in processing block 937, the set  $P_{VAT}$  of logical processes sharing virtual address translation VAT is provided to indicate that a process associated with processor  $P_i$  is sharing virtual address translation VAT; and in processing block 938, virtual address translation VAT is used to translate virtual addresses for processor  $P_i$ .

[0060] Otherwise, in processing block 934 the new virtual address translation  $VAT_i$  does not match the retrieved virtual address translation VAT and so in processing block 935 the new virtual address translation  $VAT_i$  is installed into a newly allocated entry in the TLB for processor  $P_i$ . In processing block 936, virtual address translation  $VAT_i$  is used to translate virtual addresses for processor  $P_i$ .

[0061] While a comparison of virtual address translation data may be necessary in the general case, it will be appreciated that specific implementations may permit simplifying assumptions resulting in heuristics for further optimization of the sharing of TLB entries. For example, since multiple logical processors may install different translations for the same virtual address by using corresponding page tables to drive the hardware installation of TLB entries, it may be possible to determine if a set of the logical processors are in fact using the same page tables, in which case all resulting installations of TLB entries may be shared by those processors.



[0062] One way for determining if page tables are the same may be accomplished by comparing the physical base addresses of the page tables. These base addresses, or the resulting comparisons of these base addresses, may be cached or stored in hardware to provide default sharing indications for installing virtual address translations. If the base addresses of the page tables are the same, then the resulting translations may be shared. Alternatively, if the base addresses are not the same, it does not necessarily mean that the virtual address translations may not be shared, but rather that the simplifying assumption does not apply.

[0063] Further, it may be the most probable case that the base addresses of the page tables are not changed after they are initialized. In this case, the base address comparisons may need to be performed only once. Again, if the base addresses are subsequently changed, it does not necessarily mean that the resulting translations may not be shared or even that the simplifying assumption no longer applies, but rather that the assumption may need to be reconfirmed before assigning a default sharing indication.

[0064] Figure 10 illustrates one embodiment of a computing system 1000 including a multiprocessor 1001 with a shared TLB 1002. Computing system 1000 may comprise a personal computer including but not limited to central processor 1001, graphics storage, other cache storage and local storage; system bus(es), local bus(es) and bridge(s); peripheral systems, disk and input/output systems, network systems and storage systems.

[0065] It will be appreciated that multiprocessor 1001 may comprise a single die or may comprise multiple dies. Multiprocessor 1001 may further comprise logical processors 1010 -1040, shared cache storage 1022, control logic 1004, address busses 1012, data busses 1013, bus control circuitry or other communication circuitry. Shared TLB 1002 further comprises sharing indications 1003 corresponding to virtual address translation entries in TLB 1002. When a logical processor accesses a virtual address translation entry in TLB 1002, the virtual address translation may be identified as sharable or as not sharable. A corresponding sharing indication of the sharing indications 1003 may then be provided for the virtual address translation entry.

[0066] Shared TLB 1002 supports operating-system transparent sharing of TLB entries among processors 1010-1040, which may access address spaces in common. Shared TLB 1002 further supports private TLB entries among processors 1010-1040, which for example, may each access a different physical address through identical virtual addresses. Through use of sharing indications 1003, fast and efficient virtual address translation is provided without requiring more expensive functional redundancy.

[0067] The above description is intended to illustrate preferred embodiments of the present invention. From the discussion above it should also be apparent that the invention can be modified in arrangement and detail by those skilled in the art

without departing from the principles of the present invention within the scope of the accompanying claims.

042390.P8930